

What is claimed is:

1. A method for streaming content striped in RAID 5 format from an array of disk drives to a plurality of subscribers to minimize disruptive service from a disk drive failure, said method comprising:

- 5 accessing content data on an extent-by-extent basis from a plurality of disk drives configured in an array;
- streaming the content data to the plurality of subscribers on an extent-by-extent basis, sequentially, from the plurality of disk drives;
- detecting an actual disk drive failure;
- 10 transitioning to a stream regeneration mode of operation comprising:
- reading the content data substantially simultaneously from all extents in a parity group;
- regenerating a failed portion of the content data from a
- 15 failed extent in the parity group corresponding to the failed disk drive; and
- streaming the content data in the parity group to the plurality of subscribers, extent-by-extent, immediately following the regenerating of the content data from the failed extent in the parity
- 20 group.

2. The method of claim 1, further comprising disallowing content loads upon detecting the actual disk drive failure.

- 25 3. The method of claim 1, further comprising migrating at least one subscriber to a non-failed disk drive array.

 4. The method of claim 1, further comprising migrating content to a non-failed disk drive array.

30

5. The method of claim 1, further comprising:
- sensing installment of a replacement disk drive; and
- rebuilding the content data thereon.

6. The method of claim 5, further comprising allowing content loads on the replacement disk drive.

5 7. The method of claim 6, wherein after said rebuilding step, the method further comprises:

accessing the content data on an extent-by-extent basis from the plurality of disk drives configured in the RAID 5 format; and

streaming the content data to the plurality of subscribers on an
10 extent-by-extent basis, sequentially, from the plurality of disk drives.

8. The method of claim 7, further comprising load-balancing the content data between additional disk drive arrays.

15 9. The method of claim 7, further comprising load-balancing the streams to the plurality of subscribers between additional disk drive arrays.

10. The method of claim 1, wherein the regenerating step further comprises initiating a data regeneration mode of operation comprising writing,
20 as a low priority task, recovered content data to spare extents on non-failed disk drives in the array.

11. The method of claim 10, wherein once the regenerated content data has been written to the spare extents, initiating a recovery-carousel-
25 serving mode of operation comprising streaming pseudo-sequentially, extent-by-extent, content data of each parity group to the plurality of subscribers, where the regenerated content data in a spare extent of each parity group is streamed out of sequence.

30 12. The method of claim 11, further comprising:
sensing installment of a replacement disk drive; and

writing the regenerated content data from the spare extents on the non-failed disk drives of the array to the replacement disk drive.

5 13. The method of claim 12, wherein after said writing step, the method further comprises:

 accessing content data on an extent-by-extent basis, sequentially,
 from the plurality of disk drives configured in the RAID 5 format; and
 streaming content data to the plurality of subscribers on an extent-

10 by-extent basis, sequentially, from the plurality of disk drives.

 14. The method of claim 13, further comprising allowing content loads on the replacement disk drive.

15 15. The method of claim 14, further comprising load-balancing the content data between additional disk drive arrays.

 16. The method of claim 14, further comprising load-balancing the streams to the plurality of subscribers between additional disk drive arrays.

20

 17. A method for streaming content striped in RAID 5 format from an array of disk drives to a plurality of subscribers to minimize disruptive service from a disk drive failure, said method comprising:

 accessing content data on an extent-by-extent basis from a
25 plurality of disk drives configured in an array;

 streaming the content data to the plurality of subscribers on an extent-by-extent basis, sequentially, from the plurality of disk drives;

 predicting a disk drive failure;

 writing content data from a the disk drive predicted to fail to spare
30 extents on non-failed disk drives in the array;

 detecting at least one of an actual failure and removal of the disk drive predicted to fail;

transitioning, in response to the detecting step, a recovery-carousel-serving mode of operation comprising:

5 streaming psuedo-sequentially, extent-by-extent, content data of each parity group to the plurality of subscribers, where the regenerated content data in a spare extent of each parity group is streamed out of sequence.

18. The method of claim 17, wherein the detecting step further comprises monitoring disk drive performance data selected from the group
10 consisting of a sufficiently high frequency of failed read attempts, a control signal produced by a disk failing, a thermal profile, a disk drive manufacturer detection software signal, and disk access times exceeding a predetermined threshold value.

19. The method of claim 17, wherein in an instance where the disk drive predicted to fail fails prior to said writing step, said method further comprises

transitioning to a stream regeneration mode of operation comprising:

20 reading the content data substantially simultaneously from all extents in a parity group;

 regenerating a failed portion of the content data from a failed extent in the parity group corresponding to the failed disk drive; and

25 streaming the content data in the parity group to the plurality of subscribers, extent-by-extent, immediately following the regenerating of the content data from the failed extent in the parity group.

20. The method of claim 17, further comprising migrating at least one
30 subscriber to a non-failed disk drive array.

21. The method of claim 17, further comprising migrating content to a non-failed disk drive array.

Patent 4,209,665

22. The method of claim 17, further comprising disallowing content loads upon detecting the actual failure of the disk drive predicted to fail.

5 23. The method of claim 17, further comprising:
 sensing installment of a replacement disk drive; and
 writing the regenerated content data from the spare extents
 on the non-failed disk drives of the array to the replacement disk
 drive.

10

24. The method of claim 23, wherein after said writing step, the method further comprises:

accessing content data on an extent-by-extent basis, sequentially,
from the plurality of disk drives configured in the RAID 5 format; and

15 streaming content data to the plurality of subscribers on an extent-by-
extent basis, sequentially, from the plurality of disk drives.

25. The method of claim 23, further comprising allowing content loads on the replacement disk drive.

20

26. The method of claim 24, further comprising load-balancing the content data between additional disk drive arrays.

27. The method of claim 24, further comprising load-balancing the
25 streams to the plurality of subscribers between additional disk drive arrays.